



(12) **EUROPEAN PATENT APPLICATION**

(43) Date of publication:
27.03.2019 Bulletin 2019/13

(51) Int Cl.:
G06F 12/06 ^(2006.01) **G06F 13/16** ^(2006.01)

(21) Application number: **18194678.1**

(22) Date of filing: **14.09.2018**

(84) Designated Contracting States:
**AL AT BE BG CH CY CZ DE DK EE ES FI FR GB
GR HR HU IE IS IT LI LT LU LV MC MK MT NL NO
PL PT RO RS SE SI SK SM TR**
Designated Extension States:
BA ME
Designated Validation States:
KH MA MD TN

(72) Inventors:
• **IULIANO, Luca**
Kings Langley, Hertfordshire WD4 8LZ (GB)
• **NIELD, Simon**
Kings Langley, Hertfordshire WD4 8LZ (GB)
• **ROSE, Thomas**
Kings Langley, Hertfordshire WD4 8LZ (GB)

(30) Priority: **22.09.2017 GB 201715416**

(74) Representative: **Slingsby Partners LLP**
1 Kingsway
London WC2B 6AN (GB)

(71) Applicant: **Imagination Technologies Limited**
Kings Langley,
Hertfordshire WD4 8LZ (GB)

(54) **SORTING MEMORY ADDRESS REQUESTS FOR PARALLEL MEMORY ACCESS**

(57) An apparatus configured to identify a set of M output memory addresses from a larger set of N input memory addresses containing at least one non-unique memory address, the apparatus comprising: a comparator block configured to perform comparisons of memory addresses from a set of N input memory addresses to generate a binary classification dataset that identifies a subset of addresses from the set of input addresses, where each address in the subset identified by the binary classification dataset is unique within that subset; a plurality of combination logic units, each combination logic unit being configured to: receive a predetermined selection of bits of the binary classification dataset; and sort its received predetermined selection of bits into an intermediary binary string in which the bits are ordered into a first group identifying addresses belonging to the identified subset, and a second group identifying addresses not belonging to the identified subset; and output generating logic configured to select between bits belonging to different intermediary binary strings to generate a binary output identifying a set of output memory addresses containing at least one address in the identified subset.

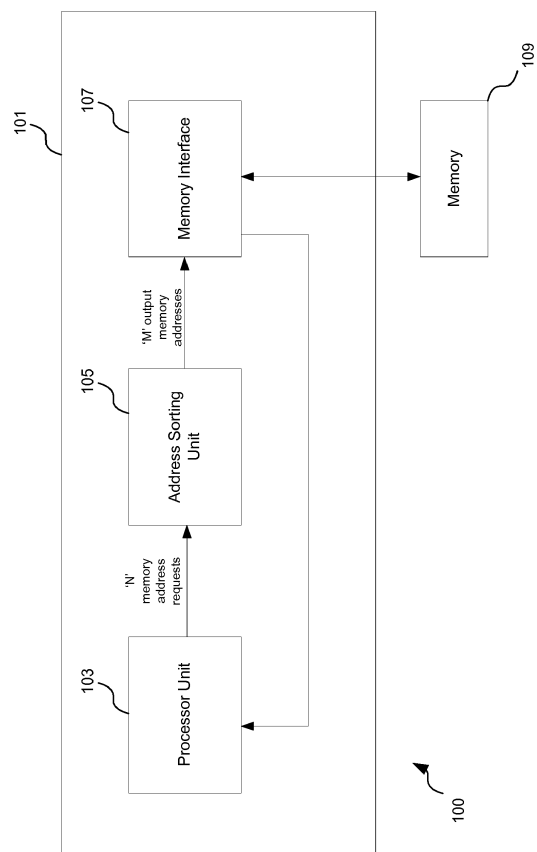


FIGURE 1

Description

FIELD

5 **[0001]** This invention relates to sorting memory address requests within a computer system for parallel memory access.

BACKGROUND

10 **[0002]** A computer processor system may include one or more processor units for performing processing operations. Each of these processor units may request access to memory (e.g. to read or write data) as part of executing instructions to perform one or more processing operations. Each memory access request may specify a memory address identifying a region of memory to be accessed. In this context, a processor unit could for example be a processor, processor core, a multi-core processor, or some other type of unit capable of executing instructions to perform one or more processing operations, such as a digital signal processor (DSP). The computer processor system could for example be a central-processor unit (CPU), a graphics processing unit (GPU).

15 **[0003]** It is common for computer processor systems to be arranged so that multiple processing operations can be performed in parallel. For example, some processor units are capable of executing multiple threads in parallel. In other examples, a computer processor system may include multiple processor units operating in parallel, each of which may execute a single thread, or multiple threads in parallel. As a consequence, a computer processor system may generate a number of memory access requests; in some systems, multiple memory access requests may be generated in a single clock cycle. In other cases, multiple memory access requests may be generated over one or more clock cycles.

20 **[0004]** To reduce the latency in the operation of the computer processor system, multiple memory accesses may be made in parallel (e.g., a specified number of memory addresses across one or more blocks of memory may be accessed in parallel). Parallelising the memory accesses may be particularly convenient when the access requests reference memory addresses within a block, or blocks, of memory not local to the processor units. For example, if the one or more processor units were implemented as part of a system-on-chip (SoC), one or more blocks of memory that can be accessed by the processor unit(s) may be located off-chip, for example to reduce the size of the chip.

25 **[0005]** The number of memory accesses that can be made in parallel may be restricted to a specified maximum value. This value may for example be limited by data bandwidth. For instance, if the processor unit(s) form part of a SoC, the rate at which data can be communicated on and off the chip may be limited by the data bandwidth limit of the memory bus used to transfer data to/from memory.

30 **[0006]** In some cases, the number of pending memory access requests may exceed the maximum number of memory accesses that can be made in parallel. Furthermore, some of the pending memory access requests may not be unique, for example the pending memory access requests may contain multiple requests to access the same memory address. Under these circumstances, a set of parallel memory access requests may contain multiple requests for the same memory address, resulting in an inefficient memory access scheme.

SUMMARY

40 **[0007]** According to the present invention there is provided an apparatus configured to identify a set of M output memory addresses from a larger set of N input memory addresses containing at least one non-unique memory address, the apparatus comprising: a comparator block configured to perform comparisons of memory addresses from a set of N input memory addresses to generate a binary classification dataset that identifies a subset of addresses from the set of input addresses, where each address in the subset identified by the binary classification dataset is unique within that subset; a plurality of combination logic units, each combination logic unit being configured to: receive a predetermined selection of bits of the binary classification dataset; and sort its received predetermined selection of bits into an intermediary binary string in which the bits are ordered into a first group identifying addresses belonging to the identified subset, and a second group identifying addresses not belonging to the identified subset; and output generating logic configured to select between bits belonging to different intermediary binary strings to generate a binary output identifying a set of output memory addresses containing at least one address in the identified subset.

45 **[0008]** The binary classification dataset may comprise N bits, each bit corresponding to a respective input memory address, the value of each bit indicating whether or not the corresponding memory address forms part of the identified subset.

50 **[0009]** The bit(s) of the binary output identifying the at least one address in the identified subset may be unordered within the output.

55 **[0010]** Each address in the set of input addresses may be associated with a validity identifier indicating whether the address is valid or invalid.

[0011] Each address in the set of input addresses that is not in the identified subset may be at least one of: an invalid

address; or equal to one of the addresses in the identified subset.

[0012] The set of input memory addresses may contain N addresses, and the comparator block may be configured to perform the comparison of input addresses by comparing address a_i with address a_j for $i = 1, \dots, N-1$ and $j = i+1, \dots, N$.

[0013] The set of input memory addresses may contain N addresses, and the comparator block may be configured to perform the comparison of input addresses by comparing address a_i with address a_j for $i = 1, \dots, N$ and $j = i-1, \dots, 1$.

[0014] The comparator block may be configured to generate the binary classification dataset by associating a first bit value with each address a_j that is determined to be both: a valid address and not equal to address a_i , and to associate a second bit value to each address a_j that is determined to be at least one of: an invalid address or equal to address a_i , the comparator block being further configured to identify an address a_j as invalid in response to determining that address a_j is equal to address a_i .

[0015] The comparator block may be configured to generate the binary classification dataset by associating a first bit value with each address a_j that is both: a valid address and not equal to any address a_i , and to associate a second bit value to each address a_j that is either: an invalid address or equal to at least one other address a_i .

[0016] The first bit values in the binary classification dataset may identify the subset of addresses in which each address is unique.

[0017] The comparator block may be further configured to generate from the comparison of input addresses a match mask indicating, for each input address, which of the other input addresses match that input address.

[0018] The apparatus may further comprise address matching logic configured to identify, using the match mask, each input address that matches the at least one address in the identified subset that is identified by the binary output.

[0019] The binary output may be an M -bit output, and the output generating logic may be configured to select between bits belonging to different intermediary binary strings to generate an M -bit binary output that identifies M addresses in the identified subset when the number of addresses in said subset is greater than or equal to M .

[0020] Each combination logic unit may be configured to sort its received predetermined selection of bits into an intermediary binary string containing fewer bits than the number of the predetermined selection of bits received by that combination logic unit.

[0021] Each combination logic unit may comprise:

a plurality of sort units each configured to:

receive a portion of the received predetermined selection of bits; and
sort the received portion of bits to group together the bits identifying addresses belonging to the identified subset to generate a preliminary binary string; and

one or more merge units, each of the one or more merge units being configured to receive a plurality of preliminary binary strings and to merge those preliminary binary strings to group together the bits from each received preliminary binary string that identify addresses belonging to the identified subset.

[0022] The output generating logic may comprise a plurality of selecting units, each configured to generate a respective bit of the binary output.

[0023] Each selecting unit may be configured to select between a pair of bits from respective intermediary binary strings to generate a bit of the binary output.

[0024] Each selecting unit may be configured to output a bit that does not identify an address in the identified subset only in response to receiving a pair of bits that both do not identify an address in the identified subset.

[0025] Each selecting unit may be associated with a pair of combination logic units and be configured to select between a pair of bits from respective intermediary binary strings formed by those pair of combination logic units.

[0026] Each combination logic unit may be configured to sort its received predetermined selection of bits into an intermediary binary string of M bits.

[0027] Each selecting unit n may be configured to select between a bit b_n from a first intermediary binary string, and a bit b_m from a second intermediary binary string, where $n = 1, \dots, M$, and $m = M+1-n$.

[0028] There may be a processor comprising the apparatus in accordance with any of the examples described herein.

[0029] The processor may be a single instruction multiple data (SIMD) processor.

[0030] In accordance with a second aspect of the present disclosure there is provided a method of identifying a set of M output memory addresses from a larger set of N input memory addresses containing at least one non-unique memory address comprising: performing a comparison of memory addresses from a set of N input memory addresses to generate a binary classification dataset that identifies a subset of addresses from the set of input addresses, where each address in the subset identified by the binary classification dataset is unique within that subset; at each of a plurality of combination logic units: receiving a predetermined selection of bits of the binary classification dataset; and sorting its received predetermined selection of bits into an intermediary binary string in which the bits are ordered into a first group

identifying addresses belonging to the identified subset, and a second group identifying addresses not belonging to the identified subset; and selecting between bits belonging to different intermediary binary strings to generate a binary output identifying a set of output memory addresses containing at least one address in the identified subset.

[0031] The binary classification dataset may comprise N bits, each bit corresponding to a respective input memory address, the value of each bit indicating whether or not the corresponding memory address forms part of the identified subset.

[0032] The bit(s) of the binary output may identify the at least one address in the identified subset are unordered within the output.

[0033] Each address in the set of input addresses may be associated with a validity identifier indicating whether the address is valid or invalid.

[0034] Each address in the set of input addresses that is not in the identified subset may be at least one of: an invalid address; or equal to one of the addresses in the identified subset.

[0035] The set of input memory addresses may contain N addresses, and the method may comprise performing the comparison of input addresses by comparing address a_i with address a_j for $i = 1, \dots, N-1$ and $j = i+1, \dots, N$.

[0036] The set of input memory addresses may contain N addresses, and the method may comprise performing the comparison of input addresses by comparing address a_i with address a_j for $i = 1, \dots, N$ and $j = i-1, \dots, 1$.

[0037] The method may comprise generating the binary classification dataset by associating a first bit value with each address a_j that is determined to be both: a valid address and not equal to address a_i , and associating a second bit value to each address a_j that is determined to be at least one of: an invalid address or equal to address a_i , the method may further comprise identifying an address a_j as invalid in response to determining that address a_j is equal to address a_i .

[0038] The method may comprise generating the binary classification dataset by associating a first bit value with each address a_i that is both: a valid address and not equal to any address a_j ; and associating a second bit value to each address a_i that is either: an invalid address or equal to at least one other address a_j .

[0039] The first bit values in the binary classification dataset may identify the subset of addresses in which each address is unique.

[0040] The method may further comprise generating from the comparison of input addresses a match mask indicating, for each input address, which of the other input addresses match that input address.

[0041] The method may further comprise identifying, using the match mask, each input address that matches the at least one address in the identified subset that is identified by the binary output.

[0042] The binary output may be an M -bit output, and the selecting step may comprise selecting between bits belonging to different intermediary binary strings to generate an M -bit binary output identifying M addresses in the identified subset when the number of addresses in said subset is greater than or equal to M .

[0043] The method may comprise, at each of the plurality of combination logic units: sorting the received predetermined selection of bits into an intermediary binary string containing fewer bits than the number of the predetermined selection of bits received by that combination logic unit.

[0044] The method may comprise, at each combination logic unit:

at each of a plurality of sort units:

receiving a portion of the received predetermined selection of bits; and
sorting the received portion of bits to group together the bits identifying addresses belonging to the identified subset to generate a preliminary binary string;

and

at each of one or more merge units:

receiving a plurality of preliminary binary strings and merging those preliminary binary strings to group together the bits from each received preliminary binary string that identify addresses belonging to the identified subset.

[0045] The selecting step may comprise, at each of a plurality of selecting units, generating a respective bit of the binary output.

[0046] The method may comprise, at each selecting unit, selecting between a pair of bits from respective intermediary binary strings to generate a bit of the binary output.

[0047] The method may comprise, at each selecting unit, outputting a bit that does not identify an address in the identified subset only in response to receiving a pair of bits that both do not identify an address in the identified subset.

[0048] The method may comprise, at each selecting unit associated with a pair of combination logic units, selecting between a pair of bits from respective intermediary binary strings formed by those pair of combination logic units.

[0049] The method may comprise, at each combination logic unit, sorting its received predetermined selection of bits into an intermediary binary string of M bits.

[0050] The method may comprise, at each selecting unit n , selecting between a bit b_n from a first intermediary binary string, and a bit b_m from a second intermediary binary string, where $n = 1, \dots, M$, and $m = M+1-n$.

[0051] There may be provided an apparatus as described herein embodied in hardware on an integrated circuit.

[0052] There may be provided a method of manufacturing, using an integrated circuit manufacturing system, an apparatus as described herein.

[0053] There may be provided an integrated circuit definition dataset that, when processed in an integrated circuit manufacturing system, configures the system to manufacture an apparatus as described herein.

[0054] There may be provided a non-transitory computer readable storage medium having stored thereon a computer readable description of an integrated circuit that, when processed in an integrated circuit manufacturing system, causes the integrated circuit manufacturing system to manufacture an apparatus as described herein.

[0055] There may be provided an integrated circuit manufacturing system configured to manufacture an apparatus as described herein.

[0056] There may be provided an integrated circuit manufacturing system comprising:

a non-transitory computer readable storage medium having stored thereon a computer readable integrated circuit description that describes an apparatus as described herein;

a layout processing system configured to process the integrated circuit description so as to generate a circuit layout description of an integrated circuit embodying the apparatus; and

an integrated circuit generation system configured to manufacture the apparatus according to the circuit layout description.

[0057] There may be provided an apparatus configured to perform any method as described herein.

[0058] There may be provided computer program code for performing a method as described herein.

[0059] There may be provided a non-transitory computer readable storage medium having stored thereon computer readable instructions that, when executed at a computer system, cause the computer system to perform a method as described herein.

BRIEF DESCRIPTION OF DRAWINGS

[0060] The present invention will now be described by way of example with reference to the accompanying drawings. In the drawings:

Figure 1 shows an overview of a computer apparatus.

Figure 2 shows an example of an address sorting unit forming part of the computer apparatus shown in figure 1.

Figure 3A illustrates the inputs received by the address sorting unit.

Figure 3B illustrates how a series of memory comparisons are performed to generate a binary classification dataset.

Figure 3C shows a match mask.

Figure 4 shows an example of a combination logic unit forming part of the address sorting unit.

Figure 5 shows an example of output generating logic forming part of the address sorting unit.

Figure 6 shows a flowchart illustrating the steps of identifying a set of output addresses from a set of input addresses.

Figure 7 shows an example computer system.

Figure 8 shows an integrated circuit manufacturing system.

DETAILED DESCRIPTION

[0061] The following description is presented by way of example to enable a person skilled in the art to make and use the invention. The present invention is not limited to the embodiments described herein and various modifications to the disclosed embodiments will be apparent to those skilled in the art. Embodiments are described by way of example only.

[0062] The present disclosure is directed to sorting a set of N input memory addresses to identify each different address (e.g. each different address value) within that set. The set of input addresses may be addresses to which access has been requested. The identified subset of different address values is then used to generate a set of M unique output memory addresses, where $M < N$. The output set of memory addresses may be such that each address within the output set is unique within that set, but may not be unique within the larger set of input addresses. That is, the input set of N memory addresses may contain a number of different address values ($\leq N$), but one or more of those address values may be included multiple times within the input set. In contrast, the term "unique", as used herein to refer to a memory address within a particular set, subset, etc., means that the memory address is not equivalent to any other memory address within that set, subset etc. If there are an insufficient number of different address values to generate an output set of addresses in which each address is unique (i.e., if the number of different address values is less than M), then the

output set of addresses contains the maximum number of different addresses.

[0063] Thus, if an input set of memory addresses contains multiple addresses with the same address value, output sets of memory addresses can be generated so that the address value is output only once. The set of output memory addresses can then be used to access the regions of memory identified by the memory addresses within the output set.

This can improve the efficiency of memory access, by reducing, or in some cases potentially eliminating, multiple accesses to the same memory address when granting access to addresses in the input set.

[0064] Examples of how the output set of memory addresses can be generated will be described in more detail below.

[0065] Figure 1 shows an example computer processing system 100. The computer processor system comprises a central processing unit (CPU) or a graphics processing unit (GPU) 101 and a memory 109. The memory 109 is shown as being external of the unit 101. If the unit 101 formed part of a system of on chip (SoC), the memory 109 could be external of the chip, i.e. memory 109 could be an off-chip memory.

[0066] The unit 101 comprises a processor unit 103, an address sorting unit 105, and a memory interface 107. The processor unit is configured to perform processing operations. To perform the processing operations, the processor unit may execute one or more instructions. The processor unit could take many different forms. For example, the processor unit could be a processor; a processor core; a multi-core processor or some other type of unit capable of executing instructions to perform one or more processing operations, such as a DSP. The processor unit may be capable of performing multiple processing operations in parallel. The processor unit could for example be a single-instruction-multiple-data (SIMD) processor. The processor unit may be capable of scheduling and executing multiple threads in parallel. Though in this example the computer processor system comprises only a single processor unit 103, in other examples the computer processor system may comprise multiple processor units. Each of these multiple processor units may operate in parallel. If the unit 101 comprises multiple processor units 103, then the units may operate independently of each other; e.g., each processor unit may be capable of independently scheduling and executing threads.

[0067] The processor unit 103 generates a set of N memory address requests. Each address request may specify a memory address value identifying a region of the memory 109 to be accessed. Each address value may be a fixed-length sequence of digits. The bit-width of each memory address value may be implementation-specific. The N address requests may be generated by the processor unit 103 in a single processor cycle, e.g. a single clock cycle. Alternatively, the N memory address requests may be generated over multiple clock cycles. In some instances, memory address requests may not be generated in discrete sets of N , but may be generated on an ad-hoc basis by the processor unit. In this case, the set of N memory address requests may refer to the N oldest pending memory requests, with the set of N address requests being updated each time a request is granted. Thus, in general, the N memory address requests may be a set of N pending requests.

[0068] The address sorting unit 105 receives the N memory addresses specified by the set of N memory access requests. The operation of the address sorting unit will be described in more detail below, but briefly: it operates to generate a set of M output memory addresses from the set of N input memory addresses specified by the received set of memory address requests. The number of memory addresses in the output set is less than the number of memory addresses in the input set (i.e., $M < N$). Each address within the set of M output memory addresses may be unique within that output set (i.e., the sorting unit outputs a unique set of M output memory addresses).

[0069] The M output memory addresses are then input into memory interface 107, which accesses the appropriate (unique) regions of memory 109 as specified by those M addresses.

[0070] Though memory 109 is shown in figure 1 as a single block, it will be appreciated that in other examples, the memory access requests may request access to multiple blocks of memory. A single set of M memory addresses may contain addresses for multiple blocks of memory.

[0071] The address sorting unit 105 is shown in more detail in figure 2. The address sorting unit 105 comprises a comparator block 201, a plurality (in this example two) of combination logic units 203A and 203B, and output generating logic 205.

[0072] The operation of the address sorting unit 105 will now be described, with reference to the flowchart in figure 6. In the examples that follow, the set of N input memory addresses received by the address sorting unit contains 32 addresses (i.e., $N = 32$), and the set of M output memory addresses contains four addresses (i.e., $M = 4$). It is to be understood that these values have been chosen merely for the purposes of illustration, and that the following description is applicable to other values of N and M .

[0073] At step 601, the address sorting unit 105 performs a comparison of the input memory addresses to generate a binary classification dataset that identifies a subset of addresses from the set of input addresses. Each address in the subset identified by the binary classification dataset is unique within that subset (but may not be unique within the input set).

[0074] Step 601 is performed by the comparator block 201. The comparator block may receive as an input the set of N input addresses. The comparator block may also receive a set of N validity identifiers, each validity identifier indicating whether a respective address in the set of input addresses is valid or invalid. Invalid addresses may arise, for example, if the number of current memory address requests is less than N . In these circumstances, the remaining input ports to the comparator block (i.e., the input ports not containing a current memory address request) may contain the address

from a previous set of input memory address requests, and be labelled as invalid. (The set of N input addresses may be received in the format of an address vector. The address vector may be formed of $N \times Add_{BW}$ bits, where Add_{BW} is the bitwidth of each memory address (and N being the number of addresses in the input set). The set of validity identifiers may be in the format of a validity vector. The validity vector may be formed of N bits, where each bit of the vector indicates whether a corresponding input address is valid or invalid. Thus, each validity identifier may be a binary identifier.

[0075] The inputs to the comparator block 201 are illustrated schematically in figure 3A. The set of N input addresses are shown generally in column 301, and the set of N validity identifiers are shown generally in column 303. It can be seen that each validity identifier corresponds to - and indicates the validity of - a respective input address. Each input memory address is a sequence of digits, in this example six digits. Each address in the input set may be of a fixed length. Each address in the input set may be of equal length. It will be understood that each address is represented as a sequence of six digits here merely for ease of illustration, and that in other implementations, each address may be formed of a different number of bits.

[0076] It is noted that, for the purposes of clarity, only seven input memory addresses and seven validity identifiers are shown in figure 3A.

[0077] The comparator block 201 performs a series of address comparisons for the input addresses to identify a subset of addresses within which each address is unique. An example scheme for comparing the input memory addresses to identify this subset will now be described with reference to figure 3B.

[0078] In the first step of the scheme, the comparator block compares the first memory address (having an address value 101101) with each of the other $N-1$ addresses in the received input set. If the addresses do not match, and the validity identifier of the address being compared to the first address indicates that the address is valid, that comparison is associated with a first binary value (in this example, the value '1'). If the addresses do match, and/or the validity identifier of the address being compared to the first address indicates the address is not valid, the comparison is associated with a second binary value (in this example, '0'). The results of the comparisons of the first address with each of the other addresses in the input set in accordance with this scheme is illustrated in column 305. It can be seen that, since address 1 matches address 6, a value of '0' is associated with that comparison. A value of '0' is also associated with the comparison of address 1 with addresses 3 and 5, because addresses 3 and 5 are indicated as invalid by their respective validity identifiers. A further step of the scheme is to identify an address as invalid in response to a determination that the address matches the first address. This may be done by changing the validity identifier associated with the address(es) that are determined to match the first address. Thus, in this example, in response to determining that address 6 matches the first address, the validity identifier associated with address 6 is altered to indicate that address 6 is not valid.

[0079] In the second step of the scheme, the second memory address is compared to each of the remaining $N-2$ addresses of the input set. As before, a first binary value ('1') is associated with a comparison if the second memory address does not match the address to which it is being compared and that address is indicated as being valid by its validity identifier. A second binary value ('0') is associated with a comparison if the second memory address matches the address to which it is being compared, and/or the validity identifier for that address indicates the address being compared to the second address is invalid. The results of the comparisons of the second memory address with the remaining $N-2$ addresses are shown in column 307. Since address 2 matches address 7, that comparison is associated with the value '0'. It will also be noticed that, though addresses 2 and 6 do not match, a value of '0' is associated with this comparison because the validity identifier for address 6 was amended to '0' to indicate the address is invalid following the prior comparison of address 6 with address 1.

[0080] The above process is repeated for each of the remaining $N-2$ addresses. Thus, for a set of N input addresses, the comparator block 201 performs a series of comparisons of address a_i with address a_j for $i = 1, \dots, N$, and $j = i+1, \dots, N$. The comparator block therefore performs $N.(N-1)/2$ address comparisons. The comparator block performs these comparisons to associate with each address a_k in the input set a binary classifier value. Put another way, the comparator block 201 may compare each address a_i with each subsequent address $a_{j=i+1 \dots N}$, repeated for each address $a_{i \in (1, \dots, N)}$ in the set of N input addresses. A first binary classification value (in this example a '1') is associated to each subsequent address a_j that is determined to be both: (i) a valid address (i.e. the address is associated with a validity identifier that indicates the address is valid); and (ii) not equal to address a_i , with the comparator block identifying said subsequent address a_j as invalid in response to determining that the address is equal to address a_i . The comparator block associates a second binary classification value (in this example, a '0') to each address a_j that is determined to be at least one of: (i) an invalid address (i.e. associated with a validity identifier that indicates the address is invalid); and (ii) equal to address a_i .

[0081] The set of binary classifier values associated with the input addresses following the operation of the comparator block 201 are indicated in figure 3B at 309. This set of binary classifier values 309 may be referred to as a binary classification dataset. In this example, the binary classification dataset is an N bit vector, where each bit of the vector contains a binary classification value for a respective input address.

[0082] The above-described scheme is a sequential scheme; i.e., the second step of the scheme is initiated after the first step of the scheme has completed, the third step of the scheme is initiated after the second step of the scheme has

completed, and so on. In general, step i is initiated only after step $i-1$ has completed. To reduce latency, the comparator block 201 may in other examples implement a parallel scheme to perform the address comparisons for the input addresses to identify a subset of addresses within which each address is unique.

[0083] The comparator block may implement the parallel scheme by performing N comparison steps. Each of the N comparison steps may be initiated in parallel, i.e. concurrently. Each comparison step i determines the binary classifier value for a respective address i in the set of N input addresses. Each step i may compare a corresponding input address i with each one of $i-1$ preceding addresses of the set of input addresses. Thus, in comparison step 1, address 1 isn't compared with any other addresses of the set of N input addresses. In comparison step N , address N is compared with each of the $N-1$ preceding addresses of the input set. Thus, each comparison step i may include $i-1$ comparisons. A total of $N(N-1)/2$ comparisons are performed in total across the N comparison steps. For each comparison step i , the comparator block compares a corresponding memory address i with a subset of $i-1$ other addresses in the input set and, if at least one of those $i-1$ subset of addresses is: 1) equal to the corresponding memory address i ; and 2) a valid address, the comparator block updates the validity identifier of corresponding memory address i to the second binary value (e.g. '0').

[0084] Put another way, the comparator block may be configured to implement the parallel scheme by performing N comparison steps, where for each comparison step i , the comparator block performs the following logical equation:

$$valid_{new,i} = valid_{old,i} \text{ AND } NOT((address_i == address_1) \text{ AND } valid_1) \text{ OR} \\ ((address_i == address_2) \text{ AND } valid_2) \text{ OR } \dots \text{ OR } ((address_i == \\ address_{i-1}) \text{ AND } valid_{i-1})) \quad (1)$$

where $valid_{new,i}$ is the value of the validity identifier for address i following comparison step i (which is the binary classifier value associated with that address); $valid_{old,i}$ is the original value of the validity identifier for address prior to comparison step i being performed; $valid_1, \dots, valid_{i-1}$ are received values of the validity identifiers for addresses 1, ..., $i-1$, and AND , NOT and OR are logical operators.

[0085] Thus, in summary, the comparator block 201 may generate the binary classification dataset by implementing a scheme in which each input address $a_{i \in \{1, \dots, N\}}$ is compared with each previous address $a_{j=1, \dots, i-1}$ in the input set to determine a binary classification value for address a_i . A first binary classification value (in this example a '1') is associated to each input address a_i that is determined to be both: (i) a valid address (i.e. the address is associated with a validity identifier that indicates the address is valid); and (ii) not equal to any previous address a_j . The comparator block associates a second binary classification value (in this example, a '0') to each input address a_i that is determined to be at least one of: (i) an invalid address (i.e. associated with a validity identifier that indicates the address is invalid); and (ii) equal to at least one previous address a_j .

[0086] Thus, in effect, each input address associated with a first binary classification value in the binary classification dataset 309 is identified as being a valid address and the first occurrence of that address within the set of input addresses when that set is serially traversed in order.

[0087] All input addresses associated with the first binary classification value in the classification dataset 309 therefore define a subset of unique addresses; i.e., each address within the subset is unique within the subset. In this example, the subset of unique addresses identified by the classification dataset includes only addresses 1, 2 and 4 for the first seven input addresses. However, each address within the subset is not necessarily unique within the set of N input addresses. For example, address 1 is unique within the subset of addresses identified by the binary classification dataset (there being no other address having a binary classification of '1' in the dataset 309 that is equal to address 1), but is not unique within the set of N input addresses (address 1 being equal to address 6). In other words, the comparator block 201 generates the binary classification dataset 309 so that, if the set of N input addresses contains multiple addresses having the same address value, only one of those addresses is associated with the first binary classification value in the classification dataset (i.e. only one of those addresses is included within the subset identified by the binary classification dataset).

[0088] The scheme described above allows all addresses in the input set to be identified that are both: valid; and the first occurrence of that address within the input set when the input set is traversed in order. In other words, each address in the set of N input addresses that is not in the subset identified by the binary classification dataset 309 (i.e. each input address associated with a '0' in the classification dataset) is at least one of: an invalid address; or equal to one of the addresses in the identified subset.

[0089] As well as generating the binary classification dataset 309, the comparator block 201 may additionally generate a match mask that indicates, for each address in the set of N input addresses, which addresses in the input set match that address. In some examples, the match mask may not take into account the validity identifiers of the input addresses. That is, the match mask may simply indicate, for each address in the set of N input addresses, which addresses in the

input set match that address independently of the validity identifiers for those addresses. In other examples, the match mask may indicate, for each address in the set of N input addresses, which addresses in the input set are both valid and match that address. The match mask may be in the form of an $N \times N$ bit matrix.

[0090] Figure 3C shows an example match mask generated by the comparator block 201 for the set of input values shown in column 301. The match mask is shown at 311. The match mask is a binary mask, with a first binary value (in this example '1') being used to denote that a pair of input addresses are both valid and match, and a second binary value (in this example '0') being used to denote that a pair of addresses do not match and/or are invalid. The match mask for each input address i may be given by a row i or column i of the matrix 311. Thus, the match mask for each input address may take the form of an N -bit vector, with each bit indicating whether address i matches a valid respective address in the input set. In other words, for each address in the input set (shown in column 301), the match mask identifies the valid addresses in the input set that match that address. The purpose of the match mask will be described in more detail below.

[0091] Retuning back to figure 6, and at step 603 each of the plurality of combination logic units 203A and 203B receives a predetermined selection of bits of the binary classification dataset. That is, each of the plurality of combination logic units receives a respective portion, or subsection, of the binary classification dataset.

[0092] The bits of the binary classification dataset may be distributed to the combination logic units so that all the bits of the dataset are allocated across the combination logic units. The bits of the binary classification dataset may be allocated to the logic units in accordance with a specified allocation scheme. This allocation scheme may be predetermined. In this particular example, where the address sorting unit 105 contains two combination logic units, each combination logic unit 203A and 203B may be configured to receive half of the bits of the binary classification dataset. For instance, logic unit 203A may receive bits 1 to $N/2$ of the binary classification dataset, and logic unit 203B may receive bits $N/2+1$ to N of the binary classification dataset. It is noted that each bit of the binary classification dataset is associated with an address value and a match mask. The following passages describe various operations performed on the binary classification bits by the combination logic units 203A, 203B and the output generating logic 205. Each of these operations is also performed on the corresponding address values and match masks so that at each stage of processing, a binary classification bit remains associated with its corresponding address value and match mask. This enables the address value and match mask to be read for each binary classification bit (i.e. each binary classifier value) output from the output generating logic 205.

[0093] At step 605, each combination logic unit 203 sorts its received selection of bits of the binary classification dataset into an intermediary binary string in which the bits are ordered into a first group identifying addresses belonging to the identified unique subset, and a second group identifying addresses not belonging to the identified unique subset. In other words, the bits of each intermediary binary string are divided into two groups: a first group of contiguous bits each having the first binary classification value, and a second group of contiguous bits each having the second binary classification value. Thus, within the intermediary binary string, the bits having the first binary classification value are contiguous, and the bits having the second binary classification value are contiguous.

[0094] Figure 4 shows a more detailed view of an exemplary combination logic unit 203. Logic unit 203A is shown here for illustration, but it will be appreciated that logic unit 203B has a similar structure. The predetermined selection of bits of the binary classification dataset received by the logic unit 203A are labelled 401. Because in this example the binary classification dataset contains 32 bits (because the set of input addresses contains 32 addresses), and there are two logic units 203A,B within the address sorting unit, each logic unit 203 receives 16 bits of the binary classification dataset. Examples values for the predetermined selection of bits 401 have been chosen to illustrate the operation of the logic unit 203A.

[0095] The logic unit 203A comprises a plurality of sort units 403A-D, and a plurality of merge units 405A-C. Each of the sort units receives a portion of the predetermined selection of bits received by the logic unit 203A. In particular, the unit 203A comprises four sort units, each configured to receive four bits of the received predetermined selection of bits 401.

[0096] Each sort unit 403A-D is configured to sort its received portion of bits to group together the bits identifying addresses belonging to the unique subset identified by the binary classification dataset to generate a respective preliminary binary string 407A-D. The sort units are also configured to order their received bits so that any bits identifying an address belonging to the unique subset form the most significant bits (MSBs) of the generated preliminary binary string. For example, the sort unit 403B receives the input string 0100. The value '1' is associated with, and thus identifies, an address belonging to the unique subset identified by the binary classification dataset. The sort unit 403B therefore operates to place the value '1' as the most significant bit of the preliminary binary string 407B.

[0097] Each merge unit 405A and 405B receives a plurality of preliminary binary strings as an input. Specifically, merge unit 405A receives as its input the two preliminary binary strings 407A and 407B generated by sort units 403A and 403B respectively; and merge unit 405B receives as its input the two preliminary binary strings 407C and 407D generated by sort units 403C and 403D respectively. Each merge unit 405A and 405B therefore receives eight bits as its input. Each merge unit 403A and 403B operates to merge its received preliminary binary strings to group together the bits from each received preliminary binary string that identify addresses belonging to the unique subset of addresses

and form a merged string. The merge units also sort the bits of the merged string so that the grouped bits identifying addresses belonging to the unique subset form the most significant bits of the merged string. The merge units further discard a number of least significant bits of the merged string to output a further preliminary binary string. The preliminary binary strings formed and output by merge units 405A and 405B are labelled 409A and 409B respectively.

[0098] In this example, the merge units 405A and 405B discard the four least significant bits of the merged string. The output of the merge units 405A and 405B is therefore a preliminary binary string of four bits. Thus, each merge unit 405A and 405B receives two preliminary binary strings as inputs, and outputs a further preliminary binary string with a number of bits equal to the number of bits of one of the received preliminary binary strings.

[0099] The merge units 405A and 405B therefore operate to generate a preliminary binary string with fewer bits than the number of bits received as an input. The operation of each of the merge units 405A and 405B can be summarised as: i) receive as an input a plurality of preliminary binary strings output by respective sort units; ii) merge the received preliminary binary strings together to generate a merged string in which bits from each received preliminary string that identify addresses belonging to the unique subset are grouped together and form the most significant bits of the merged string; and iii) discard a plurality of least significant bits of the merged string to output a further preliminary binary string with fewer bits than the number of bits received as an input.

[0100] The outputs of merge units 405A and 405B are fed as an input to a further merge unit 405C. Merge unit 405C operates to generate an intermediary binary string 411.

[0101] The merge unit 405C operates in an analogous way to merge units 405A and 405B. More specifically, merge unit 405C merges the preliminary binary strings 409A and 409B output from merge units 405A and 405B to generate a merged string in which bits from each received preliminary string 409A and 409B that identify addresses belonging to the unique subset are grouped together and form the most significant bits of the merged string; and discards a plurality of least significant bits of the merged string to generate an intermediary binary string 411.

[0102] It can be seen that the values of the intermediary binary string 411 are sorted into two groups: a first group of contiguous bits that each identify an input address belonging to the unique subset (the three bits having a value of '1'), and a second group containing (in this example) a single bit that identifies an input address not belonging to the unique subset (the bit having a value of '0'). The bits of the intermediary binary string 411 are sorted, or ordered, so that the bits of the first group form the most significant bits of the intermediary binary output.

[0103] Merge unit 405C receives two preliminary four-bit binary strings 409A and 409B, and generates an intermediary binary string with four bits. Thus, merge unit 405C receives as its input two preliminary binary strings output from respective merge units 405A and 405B, and generates an intermediary binary string with a number of bits equal to the number of bits of one of the received preliminary binary strings. The intermediary binary string 411 generated by the combination logic unit 203A therefore contains fewer bits than the number of bits 401 of the binary classification dataset received at the logic unit as an input.

[0104] Combination logic unit 203B operates in an analogous manner to logic unit 203A to generate a second intermediary binary string.

[0105] The two intermediary binary strings generated by the logic units 203A and 203B are input into the output generating logic 205.

[0106] At step 607, the output generating logic 205 selects between bits belonging to the two different intermediary binary strings to generate a binary output identifying a set of memory addresses containing at least one address in the unique subset. The set of memory addresses identified by the binary output may be referred to as an output set of memory addresses.

[0107] Figure 5 shows a more detailed view of exemplary output generating logic 205, and illustrates how the output generating logic may select between bits belonging to different intermediary binary strings output from respective combination logic units to generate the binary output 503. The binary output is an M bit output. It may be in the format of an M bit vector. As mentioned above, each bit of the binary output 503 is associated with an address value from the set of N input addresses, and a match mask indicating which addresses in the input set are valid and match that address value.

[0108] The logic 205 receives as its input the intermediary binary string 411 output from logic unit 203A and the intermediary binary string 413 output from logic unit 203B. Example values of the intermediary binary strings have been chosen to illustrate the operating principles of the output generating logic 205.

[0109] The output generating logic 205 comprises a plurality of selecting units 501A-D that each generate a respective bit of the binary output 503. Each selecting unit 501A-D could be a multiplexer ("mux"), for example. Since the binary output is an M bit output (where $M = 4$ in this example), the output generating logic may contain M selecting units.

[0110] Each selecting unit 501A-D is configured to receive a pair of bits from respective intermediary binary strings 411 and 413; that is, each bit in a received pair is from a different intermediary binary string. The selecting unit operates to select one of the received pair of bits as a bit of the binary output 503. Each selecting unit may operate as an OR gate. That is, if one of the input pair of bits is '1', the selecting unit outputs that bit. If both of the input bits are '1', the selecting unit selects one of the bits, either at random, or the selecting unit may be configured to select the bit output from a specified combination logic unit.

[0111] The selection units are arranged so that one selecting unit (in this example unit 501A) is configured to receive the most significant bit of the first intermediary binary string 411 and the least significant bit of the second intermediary binary string 413; one of the selecting units (in this example unit 501B) is configured to receive the second MSB of the intermediary string 411 and the second LSB of the second intermediary string 413; one selecting unit (in this example unit 501C) is configured to receive the third MSB of the intermediary string 411 and the third LSB of the intermediary string 413; and one selecting unit (in this example unit 501D) is configured to receive the fourth MSB of the intermediary binary string 411 (which in this example is the LSB of the string 411) and the fourth LSB of the binary string 413 (which in this example is the MSB of string 413).

[0112] This arrangement is advantageous because, coupled with the ordered intermediary strings output by the combination logic units, it enables the logic 205 to maximise the amount of bits in the binary output 503 that identify addresses within the unique subset for a given set of input intermediary binary strings. For example, if instead the outputs of the logic units 203A and 203B were paired so that selecting unit 501A received the MSB of both intermediary strings; selecting unit 501B received the second MSB of both intermediary strings etc., then selecting unit 501D would receive a pair of '0' bits as its inputs, meaning the binary output 503 would only contain three '1' bits, instead of four as in the illustrated arrangement.

[0113] Which selecting unit receives which combination of outputs from the combination logic units can be varied whilst still maximising the amount of bits in the binary output 503 that identify addresses within the unique subset for a given set of input intermediary binary strings. Thus, in general, the output generating logic 205 may be arranged so that a selecting unit n is configured to select between a bit b_n from a first intermediary binary string, and a bit b_m from a second intermediary binary string, where $n = 1 \dots M$, and $m = M + 1 - n$. It is noted here that under this notation, b_1 may refer to the most significant bit of an intermediary binary string, and b_M to the least significant bit of the intermediary binary string. It is also noted that under this notation the labelling of the selecting units may be arbitrary. That is, there need not be any correlation between selecting unit n and the bit number of the binary output 503, e.g. selecting unit 1 need not output the MSB of the binary output.

[0114] The configuration of the output generating logic 205 means that the bits of the binary output 503 are not ordered (in contrast to, for example, the intermediary strings output by the combination logic units). That is, bit values of '1' and '0' need not be grouped or sorted in any particular order in the binary output 503. In this regard, the operation of the output generating logic 205 can be contrasted with the merge units 405A-C. Each of the merge units and the output generating logic 205 operate to output four bits from a received input of eight bits. However, the output generating logic 205 does not sort, or order, the output bits to group certain bits together, unlike the merge units, which group together bits having a bit value of '1' in their outputs (as illustrated in figure 4). The output generating logic 205 may be said to generate a binary output with a loose packing of bits. This is advantageous because it reduces the amount of hardware (e.g. the number of gates and/or MUXs) needed to generate the binary output from the input intermediary strings 411 and 413. Thus, the output generating logic may be implemented in a more hardware-efficient manner than the merge units, reducing the overall hardware requirements of the address sorting unit 105.

[0115] The binary output 503 and the associated memory addresses of the unique subset identified by that binary output are input into the memory interface 107. The memory interface can then access regions of the memory 109 specified by those memory addresses.

[0116] Once the memory access to those regions of memory has completed, the memory interface 107 can use information from the match mask 311 generated by the comparator block to determine which of the N input memory access requests received by the sorting unit 105 have been completed (e.g. using memory matching logic). For example, it can be seen from the exemplary match mask 311 that accessing the memory address 101101 means both memory request 1 and memory request 6 are completed. The address sorting unit may output the match mask directly to the memory interface. Alternatively, the address sorting unit may further comprise an address matching logic that generates $M N$ -bit masks for each binary output, each N bit mask corresponding to a respective bit of the binary output and indicating which addresses in the set of N input addresses match the output address identified by that bit of the binary output. The address matching logic may be coupled to the comparator block 201 and the output generating logic. The address matching logic may be configured to generate the M masks from the match mask generated by the comparator block 201, and the binary output generated by the output generating logic. The address sorting unit may be configured to output the $M N$ -bit masks, for example to the memory interface unit to determine which of the input memory address requests have been completed.

[0117] The above-described approach to generating the binary output 503 means that, if the binary classification dataset generated by the comparator block identifies M or more input addresses as belonging to the unique subset, the combination logic units and output generating logic will operate to generate a binary output 503 that identifies M addresses of that subset. This is advantageous because it means the number of addresses belonging to the unique subset identified by a single binary output is maximised.

[0118] It will be appreciated that in some cases the binary classification dataset may identify a number of addresses belonging to the unique subset that is greater than M . In this case, not all the addresses in the subset can be identified

in a single binary output. Multiple cycles may therefore need to be performed to generate multiple binary outputs from the same set of N input memory addresses. Whether another cycle is to be performed to generate a further binary output using the same set of input memory addresses may be controlled by an output flag. The output flag may be enabled when it is determined that not all addresses in the unique subset have yet been identified by the binary output(s), and indicates that a further cycle to generate a new binary output is required using the same set of N input addresses. The output flag may be disabled when it is determined that all the addresses in the unique subset have been identified by the generated binary output(s), and indicates that a new set of N input addresses can be received by the address sorting unit 105.

[0119] This output flag may be generated by a monitoring unit forming part of the address sorting unit 105. The monitoring unit may be coupled to both the output generating logic and the comparator block. The monitoring unit may determine the number of input addresses in the unique subset from the binary classification dataset, and may monitor the number of those addresses that have been identified by binary outputs generated by the output generating logic. From these parameters, the monitoring unit may generate either a set or non-set output flag that is input into the comparator block.

[0120] The above examples describe an address sorting unit with two combination logic units. It will be appreciated that in other examples, the address sorting unit may contain more than two combination logic units. Furthermore, though each combination logic unit is shown as comprising a tree-network structured arrangement of four sort units, two merge units configured to receive the outputs of the sort units, and a single other merge unit configured to receive the output of the two merge units, it will be appreciated that different numbers and arrangements of sort and merge units are possible. The arrangements described herein are convenient when the set of input addresses contains 32 addresses and the set of output addresses contains 4 addresses, but it will be appreciated that other arrangements may be more convenient for differently-sized input and output address sets.

[0121] In the examples above, the comparator block 201 is described as operating to compare input addresses a_i and a_j (e.g. using a sequential or parallel scheme). It will be appreciated that the numbers used to label each address (e.g. address '1', address '2' etc.) referenced by may be arbitrarily assigned. That is, the above-described examples may not require a specific number label be assigned to each address, so long as each address within the input set N is labelled uniquely. Moreover, the number labels assigned to each address need not necessarily imply any limitation on the position of that address within the address vector received by the comparator block. In some examples, address '1' may be used to refer to the address occupying the most significant bits of the received address vector, and address 'N' used to refer to the address occupying the least significant bits of the received address vector; in other examples, address '1' may be used to refer to the address occupying the least significant bits of the received address vector, and address 'N' used to refer to the address occupying the most significant bits of the received address vector. Other approaches to labelling the addresses within the address vector are also possible.

[0122] The above examples illustrate how an address sorting unit can operate to identify a set of M unique addresses from a larger input set of N addresses, where the input set contains at least one non-unique address. However, it will be appreciated that the examples described herein will equally operate to identify a set of M unique addresses from an input set of N unique addresses. Thus, the set of N input addresses may be unique (i.e. contain no non-unique addresses), or contain at least one non-unique address.

[0123] Figure 7 shows a computer system in which the graphics processing systems described herein may be implemented. The computer system comprises a CPU 702, a GPU 704, a memory 706 and other devices 714, such as a display 716, speakers 718 and a camera 708. A processing block 710 (e.g. corresponding to address sorting unit 105) is shown implemented on the GPU 704. The components of the computer system can communicate with each other via a communications bus 720. Either or both of CPU 702 and GPU 704 could be SIMD processors. In other examples, the processing block 710 may be implemented on the CPU 702. A store 712 (e.g. corresponding to memory 109) is implemented as part of the memory 706.

[0124] The address sorting unit, combination logic unit and output generating logic of Figures 2, 4 and 5 are shown as comprising a number of functional blocks. This is schematic only and is not intended to define a strict division between different logic elements of such entities. Each functional block may be provided in any suitable manner. It is to be understood that intermediate values described herein as being formed by a unit, logic or block need not be physically generated by the unit/logic/block at any point and may merely represent logical values which conveniently describe the processing performed by the unit/logic/block between its input and output.

[0125] The apparatuses described herein (e.g. the address sorting unit 105 or processing unit 101) may be embodied in hardware on an integrated circuit. The apparatuses described herein may be configured to perform any of the methods described herein. Generally, any of the functions, methods, techniques or components described above can be implemented in software, firmware, hardware (e.g., fixed logic circuitry), or any combination thereof. The terms "module," "functionality," "component," "element," "unit," "block" and "logic" may be used herein to generally represent software, firmware, hardware, or any combination thereof. In the case of a software implementation, the module, functionality, component, element, unit, block or logic represents program code that performs the specified tasks when executed on

a processor. The algorithms and methods described herein could be performed by one or more processors executing code that causes the processor(s) to perform the algorithms/methods. Examples of a computer-readable storage medium include a random-access memory (RAM), read-only memory (ROM), an optical disc, flash memory, hard disk memory, and other memory devices that may use magnetic, optical, and other techniques to store instructions or other data and that can be accessed by a machine.

[0126] The terms computer program code and computer readable instructions as used herein refer to any kind of executable code for processors, including code expressed in a machine language, an interpreted language or a scripting language. Executable code includes binary code, machine code, bytecode, code defining an integrated circuit (such as a hardware description language or netlist), and code expressed in a programming language code such as C, Java or OpenCL. Executable code may be, for example, any kind of software, firmware, script, module or library which, when suitably executed, processed, interpreted, compiled, executed at a virtual machine or other software environment, cause a processor of the computer system at which the executable code is supported to perform the tasks specified by the code.

[0127] A processor, computer, or computer system may be any kind of device, machine or dedicated circuit, or collection or portion thereof, with processing capability such that it can execute instructions. A processor may be any kind of general purpose or dedicated processor, such as a CPU, GPU, System-on-chip, state machine, media processor, an application-specific integrated circuit (ASIC), a programmable logic array, a field-programmable gate array (FPGA), or the like. A computer or computer system may comprise one or more processors. It is also intended to encompass software which defines a configuration of hardware as described herein, such as HDL (hardware description language) software, as is used for designing integrated circuits, or for configuring programmable chips, to carry out desired functions. That is, there may be provided a computer readable storage medium having encoded thereon computer readable program code in the form of an integrated circuit definition dataset that when processed in an integrated circuit manufacturing system configures the system to manufacture an apparatus configured to perform any of the methods described herein, or to manufacture an apparatus comprising any apparatus described herein. An integrated circuit definition dataset may be, for example, an integrated circuit description.

[0128] There may be provided a method of manufacturing, at an integrated circuit manufacturing system, an apparatus as described herein. There may be provided an integrated circuit definition dataset that, when processed in an integrated circuit manufacturing system, causes the method of manufacturing an apparatus to be performed.

[0129] An integrated circuit definition dataset may be in the form of computer code, for example as a netlist, code for configuring a programmable chip, as a hardware description language defining an integrated circuit at any level, including as register transfer level (RTL) code, as high-level circuit representations such as Verilog or VHDL, and as low-level circuit representations such as OASIS (RTM) and GDSII. Higher level representations which logically define an integrated circuit (such as RTL) may be processed at a computer system configured for generating a manufacturing definition of an integrated circuit in the context of a software environment comprising definitions of circuit elements and rules for combining those elements in order to generate the manufacturing definition of an integrated circuit so defined by the representation. As is typically the case with software executing at a computer system so as to define a machine, one or more intermediate user steps (e.g. providing commands, variables etc.) may be required in order for a computer system configured for generating a manufacturing definition of an integrated circuit to execute code defining an integrated circuit so as to generate the manufacturing definition of that integrated circuit.

[0130] An example of processing an integrated circuit definition dataset at an integrated circuit manufacturing system so as to configure the system to manufacture an apparatus (e.g. an address sorting unit 105 or processing unit 101) will now be described with respect to Figure 8.

[0131] Figure 8 shows an example of an integrated circuit (IC) manufacturing system 802 which is configured to manufacture an apparatus as described in any of the examples herein. In particular, the IC manufacturing system 802 comprises a layout processing system 804 and an integrated circuit generation system 806. The IC manufacturing system 802 is configured to receive an IC definition dataset (e.g. defining an apparatus as described in any of the examples herein), process the IC definition dataset, and generate an IC according to the IC definition dataset (e.g. which embodies an apparatus as described in any of the examples herein). The processing of the IC definition dataset configures the IC manufacturing system 802 to manufacture an integrated circuit embodying an apparatus as described in any of the examples herein.

[0132] The layout processing system 804 is configured to receive and process the IC definition dataset to determine a circuit layout. Methods of determining a circuit layout from an IC definition dataset are known in the art, and for example may involve synthesising RTL code to determine a gate level representation of a circuit to be generated, e.g. in terms of logical components (e.g. NAND, NOR, AND, OR, MUX and FLIP-FLOP components). A circuit layout can be determined from the gate level representation of the circuit by determining positional information for the logical components. This may be done automatically or with user involvement in order to optimise the circuit layout. When the layout processing system 804 has determined the circuit layout it may output a circuit layout definition to the IC generation system 806. A circuit layout definition may be, for example, a circuit layout description.

[0133] The IC generation system 806 generates an IC according to the circuit layout definition, as is known in the art.

For example, the IC generation system 806 may implement a semiconductor device fabrication process to generate the IC, which may involve a multiple-step sequence of photo lithographic and chemical processing steps during which electronic circuits are gradually created on a wafer made of semiconducting material. The circuit layout definition may be in the form of a mask which can be used in a lithographic process for generating an IC according to the circuit definition.

Alternatively, the circuit layout definition provided to the IC generation system 806 may be in the form of computer-readable code which the IC generation system 806 can use to form a suitable mask for use in generating an IC.

[0134] The different processes performed by the IC manufacturing system 802 may be implemented all in one location, e.g. by one party. Alternatively, the IC manufacturing system 802 may be a distributed system such that some of the processes may be performed at different locations, and may be performed by different parties. For example, some of the stages of: (i) synthesising RTL code representing the IC definition dataset to form a gate level representation of a circuit to be generated, (ii) generating a circuit layout based on the gate level representation, (iii) forming a mask in accordance with the circuit layout, and (iv) fabricating an integrated circuit using the mask, may be performed in different locations and/or by different parties.

[0135] In other examples, processing of the integrated circuit definition dataset at an integrated circuit manufacturing system may configure the system to manufacture an apparatus without the IC definition dataset being processed so as to determine a circuit layout. For instance, an integrated circuit definition dataset may define the configuration of a reconfigurable processor, such as an FPGA, and the processing of that dataset may configure an IC manufacturing system to generate a reconfigurable processor having that defined configuration (e.g. by loading configuration data to the FPGA).

[0136] In some embodiments, an integrated circuit manufacturing definition dataset, when processed in an integrated circuit manufacturing system, may cause an integrated circuit manufacturing system to generate a device as described herein. For example, the configuration of an integrated circuit manufacturing system in the manner described above with respect to Figure 8 by an integrated circuit manufacturing definition dataset may cause a device as described herein to be manufactured.

[0137] In some examples, an integrated circuit definition dataset could include software which runs on hardware defined at the dataset or in combination with hardware defined at the dataset. In the example shown in Figure 8, the IC generation system may further be configured by an integrated circuit definition dataset to, on manufacturing an integrated circuit, load firmware onto that integrated circuit in accordance with program code defined at the integrated circuit definition dataset or otherwise provide program code with the integrated circuit for use with the integrated circuit.

[0138] The graphics processing systems described herein may be embodied in hardware on an integrated circuit. The graphics processing systems described herein may be configured to perform any of the methods described herein.

[0139] The implementation of concepts set forth in this application in devices, apparatus, modules, and/or systems (as well as in methods implemented herein) may give rise to performance improvements when compared with known implementations. The performance improvements may include one or more of increased computational performance, reduced latency, increased throughput, and/or reduced power consumption. During manufacture of such devices, apparatus, modules, and systems (e.g. in integrated circuits) performance improvements can be traded-off against the physical implementation, thereby improving the method of manufacture. For example, a performance improvement may be traded against layout area, thereby matching the performance of a known implementation but using less silicon. This may be done, for example, by reusing functional blocks in a serialised fashion or sharing functional blocks between elements of the devices, apparatus, modules and/or systems. Conversely, concepts set forth in this application that give rise to improvements in the physical implementation of the devices, apparatus, modules, and systems (such as reduced silicon area) may be traded for improved performance. This may be done, for example, by manufacturing multiple instances of a module within a predefined area budget.

[0140] The applicant hereby discloses in isolation each individual feature described herein and any combination of two or more such features, to the extent that such features or combinations are capable of being carried out based on the present specification as a whole in the light of the common general knowledge of a person skilled in the art, irrespective of whether such features or combinations of features solve any problems disclosed herein, and without limitation to the scope of the claims. The applicant indicates that aspects of the present invention may consist of any such individual feature or combination of features. In view of the foregoing description it will be evident to a person skilled in the art that various modifications may be made within the scope of the invention.

Claims

1. An apparatus configured to identify a set of M output memory addresses from a larger set of N input memory addresses containing at least one non-unique memory address, the apparatus comprising:

a comparator block configured to perform comparisons of memory addresses from a set of N input memory

addresses to generate a binary classification dataset that identifies a subset of addresses from the set of input addresses, where each address in the subset identified by the binary classification dataset is unique within that subset;

a plurality of combination logic units, each combination logic unit being configured to: receive a predetermined selection of bits of the binary classification dataset; and sort its received predetermined selection of bits into an intermediary binary string in which the bits are ordered into a first group identifying addresses belonging to the identified subset, and a second group identifying addresses not belonging to the identified subset; and

output generating logic configured to select between bits belonging to different intermediary binary strings to generate a binary output identifying a set of output memory addresses containing at least one address in the identified subset.

2. An apparatus as claimed in claim 1, wherein the binary classification dataset comprises N bits, each bit corresponding to a respective input memory address, the value of each bit indicating whether or not the corresponding memory address forms part of the identified subset.

3. An apparatus as claimed in claim 1 or 2, wherein the bit(s) of the binary output identifying the at least one address in the identified subset are unordered within the output.

4. An apparatus as claimed in any preceding claim, wherein the set of input memory addresses contains N addresses, and the comparator block is configured to perform the comparison of input addresses by comparing each address a_i in the set of input memory addresses with each subsequent address a_j in the set of input memory addresses; or perform the comparison of input addresses by comparing each address a_i in the set of input memory addresses with each previous address a_j in the set of input memory address.

5. An apparatus as claimed in claim 4, wherein each address in the set of input addresses is associated with a validity identifier indicating whether the address is valid or invalid and the comparator block is configured to generate the binary classification dataset by associating a first bit value with each address a_i that is determined to be both: a valid address and not equal to address a_j , and to associate a second bit value to each address a_j that is determined to be at least one of: an invalid address or equal to address a_i , the comparator block being further configured to identify an address a_j as invalid in response to determining that address a_j is equal to address a_i .

6. An apparatus as claimed in any preceding claim, wherein the comparator block is further configured to generate from the comparison of input addresses a match mask indicating, for each input address, which of the other input addresses match that input address, the apparatus further comprising address matching logic configured to identify, using the match mask, each input address that matches the at least one address in the identified subset that is identified by the binary output.

7. An apparatus as claimed in any preceding claim, wherein the binary output is an M -bit output, and the output generating logic is configured to select between bits belonging to different intermediary binary strings to generate an M -bit binary output that identifies M addresses in the identified subset when the number of addresses in said subset is greater than or equal to M .

8. An apparatus as claimed in any preceding claim, wherein each combination logic unit comprises:

a plurality of sort units each configured to:

receive a portion of the received predetermined selection of bits; and

sort the received portion of bits to group together the bits identifying addresses belonging to the identified subset to generate a preliminary binary string; and

one or more merge units, each of the one or more merge units being configured to receive a plurality of preliminary binary strings and to merge those preliminary binary strings to group together the bits from each received preliminary binary string that identify addresses belonging to the identified subset.

9. An apparatus as claimed in any preceding claim wherein the output generating logic comprises a plurality of selecting units, each configured to generate a respective bit of the binary output and each selecting unit is configured to select between a pair of bits from respective intermediary binary strings to generate a bit of the binary output.

10. An apparatus as claimed in claim 9, wherein each selecting unit is configured to output a bit that does not identify an address in the identified subset only in response to receiving a pair of bits that both do not identify an address in the identified subset.

11. An apparatus as claimed in claim 9 or 10, wherein each selecting unit is associated with a pair of combination logic units and is configured to select between a pair of bits from respective intermediary binary strings formed by those pair of combination logic units.

12. An apparatus as claimed in any preceding claim, wherein each combination logic unit is configured to sort its received predetermined selection of bits into an intermediary binary string of M bits, wherein each selecting unit n is configured to select between a bit b_n from a first intermediary binary string, and a bit b_m from a second intermediary binary string, where $n = 1, \dots, M$, and $m = M+1-n$.

13. A processor comprising the apparatus of any preceding claim.

14. A method of identifying a set of M output memory addresses from a larger set of N input memory addresses containing at least one non-unique memory address comprising:

performing a comparison of memory addresses from a set of N input memory addresses to generate a binary classification dataset that identifies a subset of addresses from the set of input addresses, where each address in the subset identified by the binary classification dataset is unique within that subset;

at each of a plurality of combination logic units: receiving a predetermined selection of bits of the binary classification dataset; and sorting its received predetermined selection of bits into an intermediary binary string in which the bits are ordered into a first group identifying addresses belonging to the identified subset, and a second group identifying addresses not belonging to the identified subset; and

selecting between bits belonging to different intermediary binary strings to generate a binary output identifying a set of output memory addresses containing at least one address in the identified subset.

15. A computer readable storage medium having stored thereon an integrated circuit definition dataset that, when processed in an integrated circuit manufacturing system, configures the system to manufacture an apparatus as claimed in any of claims 1 to 12.

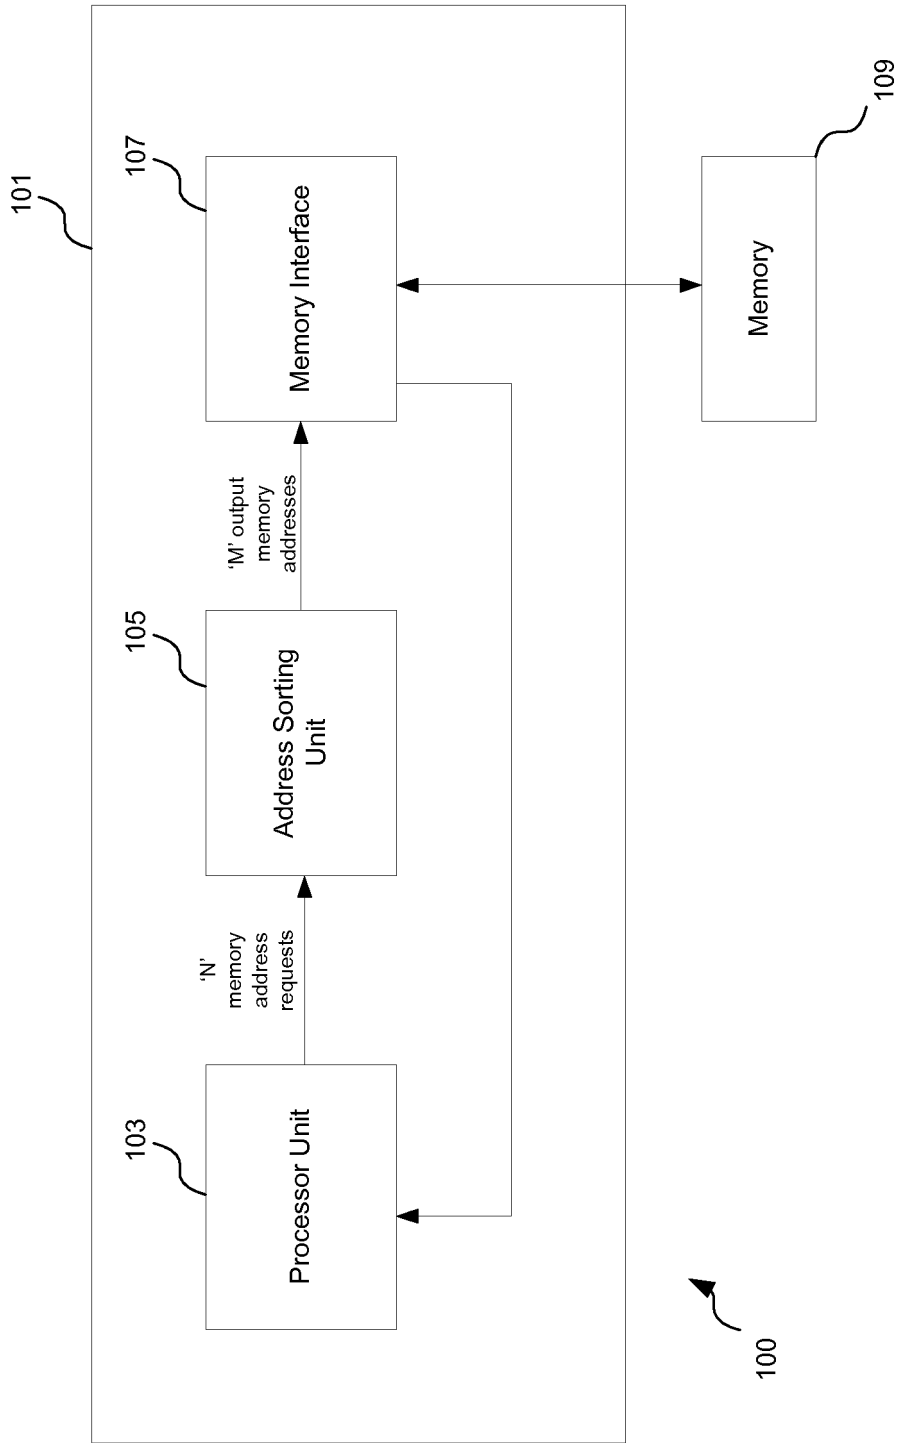


FIGURE 1

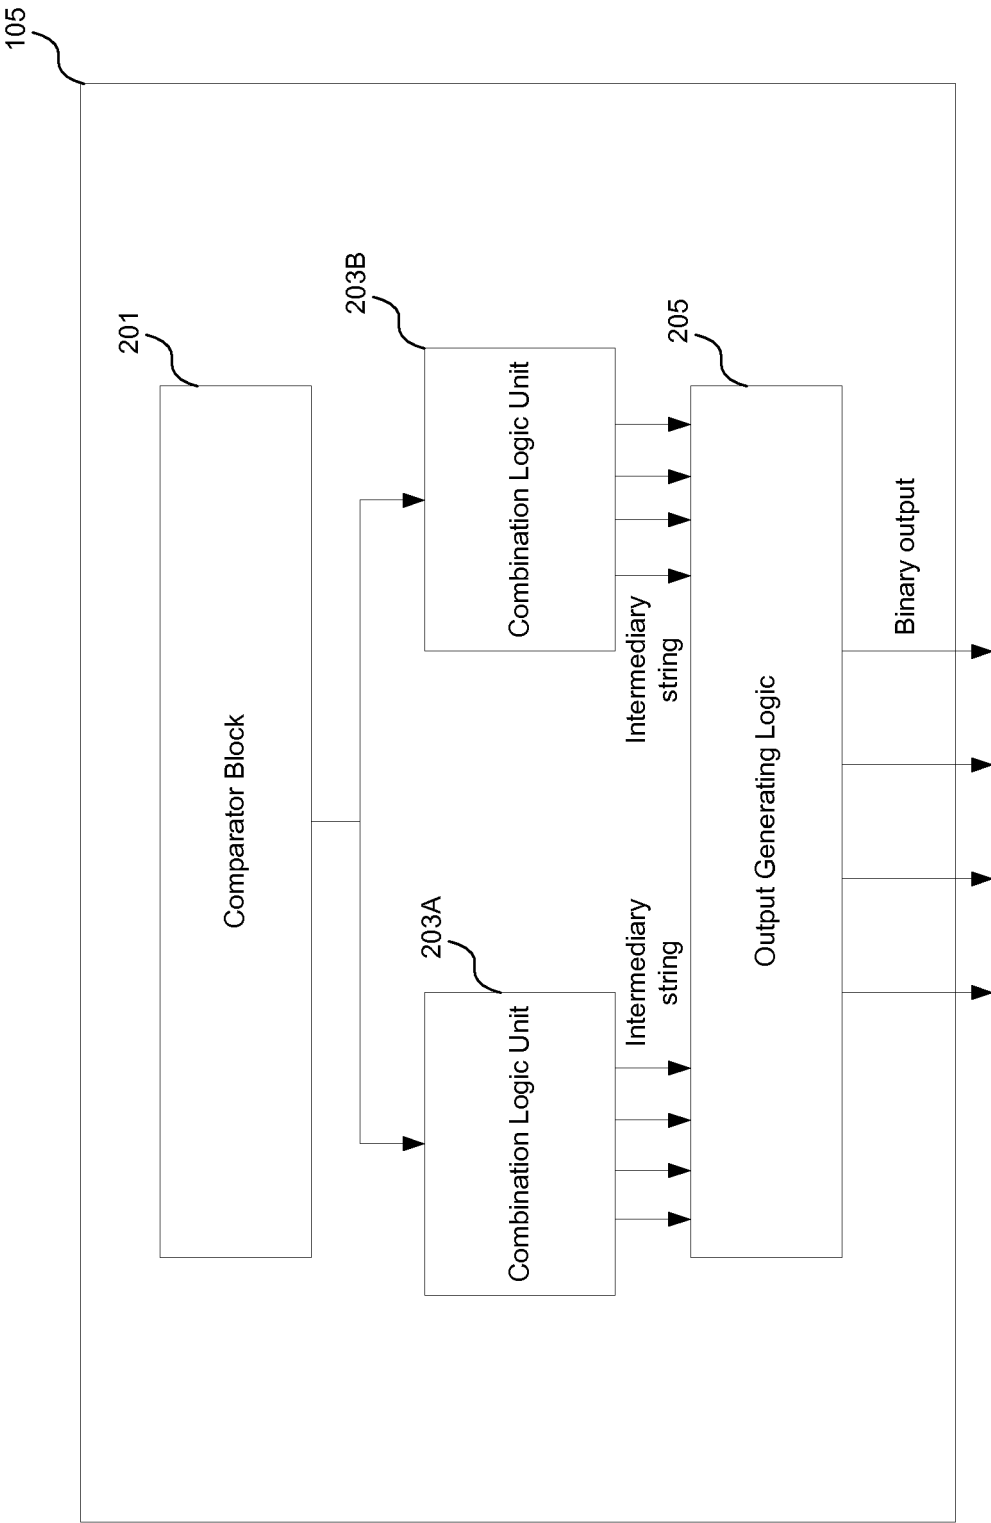


FIGURE 2

	Validity Identifier	Address
1	1	101101
2	1	011101
3	0	111000
4	1	010100
5	0	101010
6	1	101101
7	1	011101
	⋮	⋮
N	⋮	⋮

303 301

FIGURE 3A

	Validity Identifier	Address	Binary Classifier							
1	1	101101	1							
2	1	011101	1	1						
3	0	111000	0	0	0					
4	1	010100	1	1	1	1				
5	0	101010	0	0	0	0	0			
6	1 0	101101	0	0	0	0	0	0		
7	1 0	011101	1	0	0	0	0	0	0	
	⋮	⋮								⋮
N	⋮	⋮								

303 301 305 307 309

FIGURE 3B

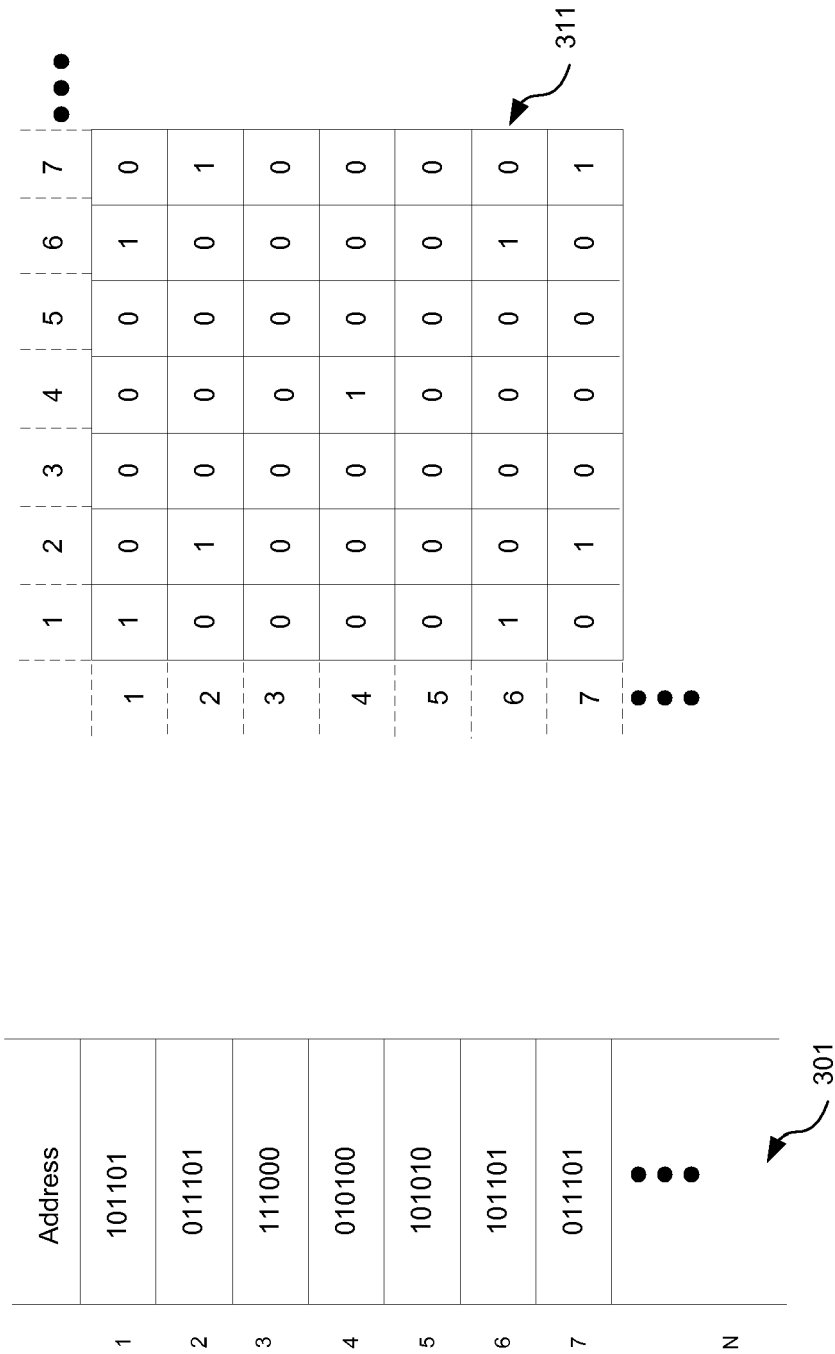


FIGURE 3C

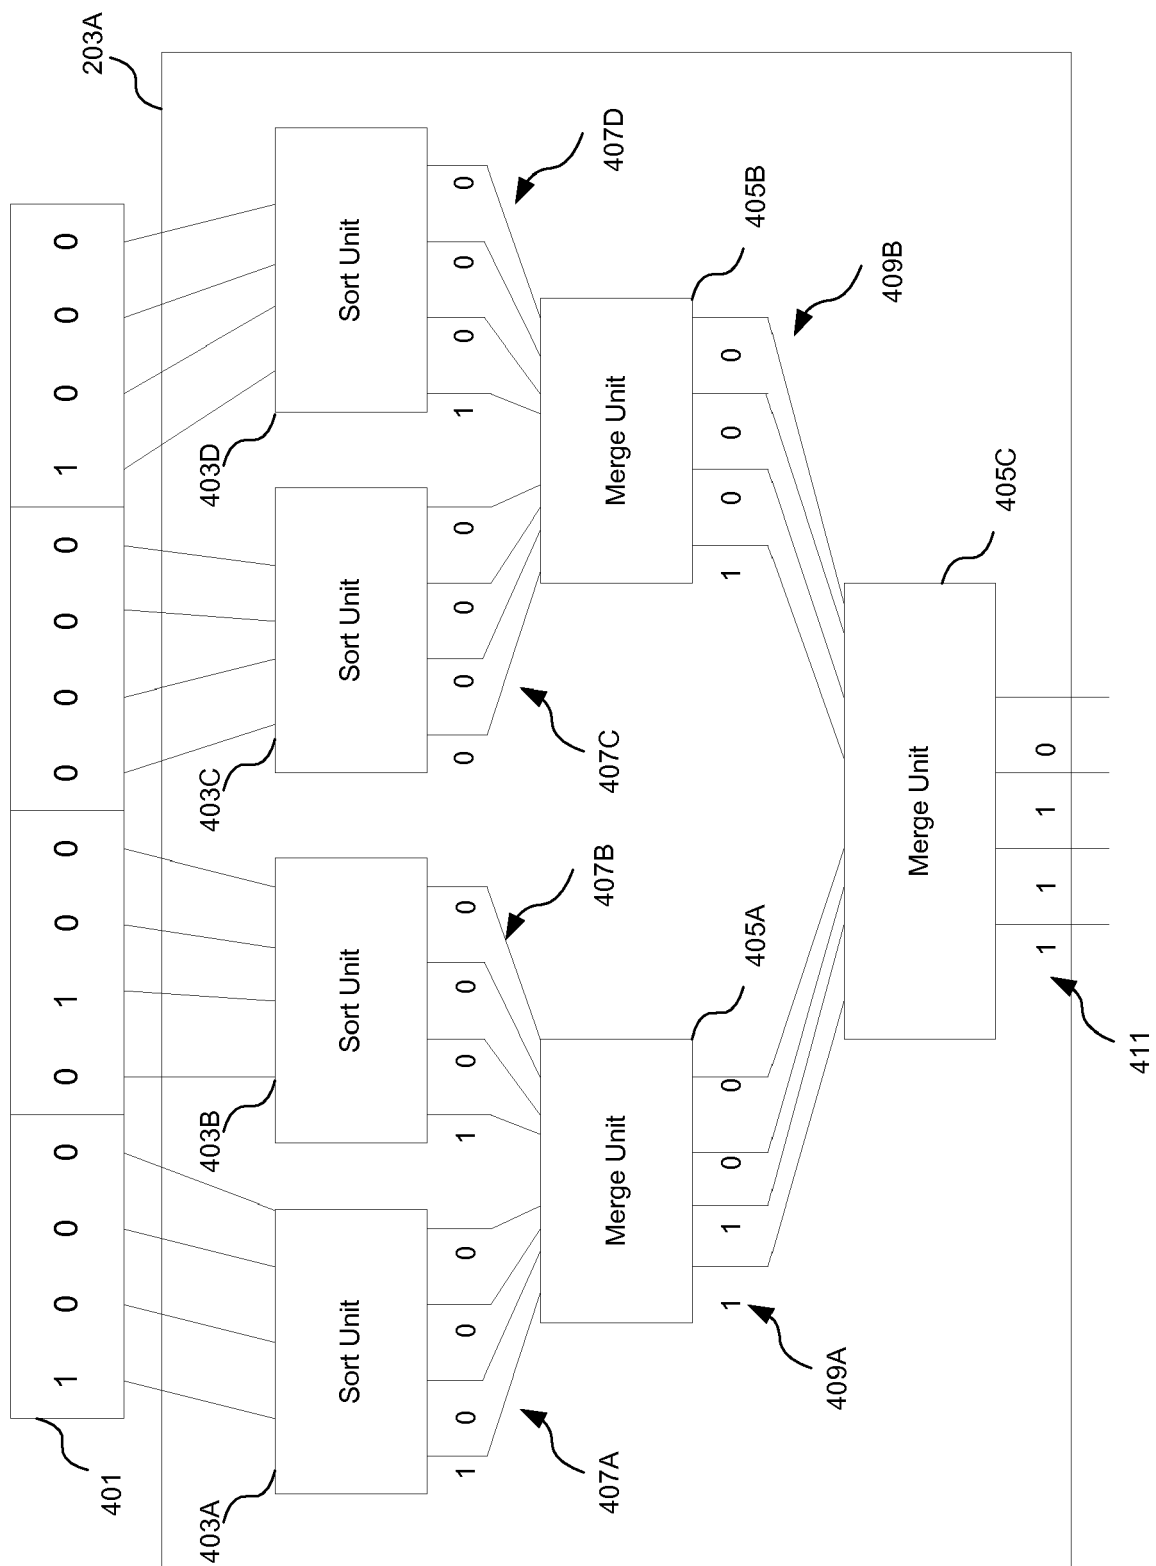


FIGURE 4

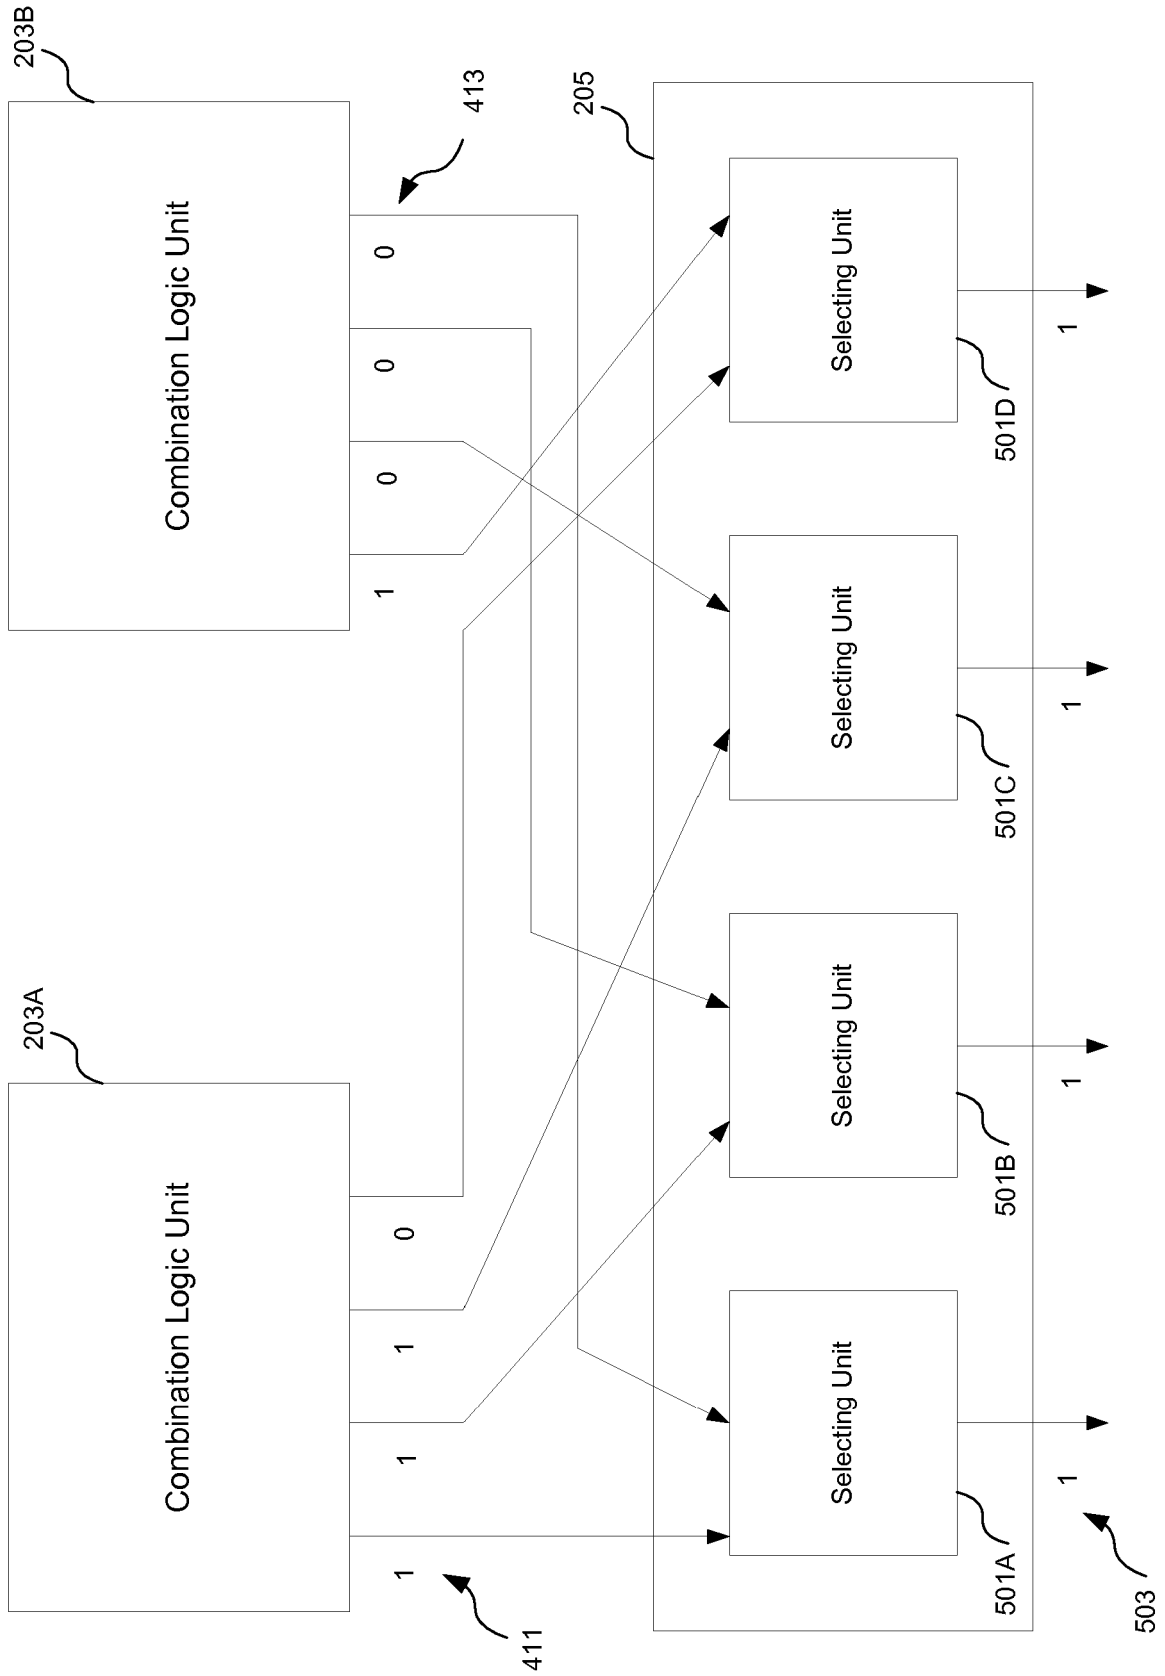


FIGURE 5

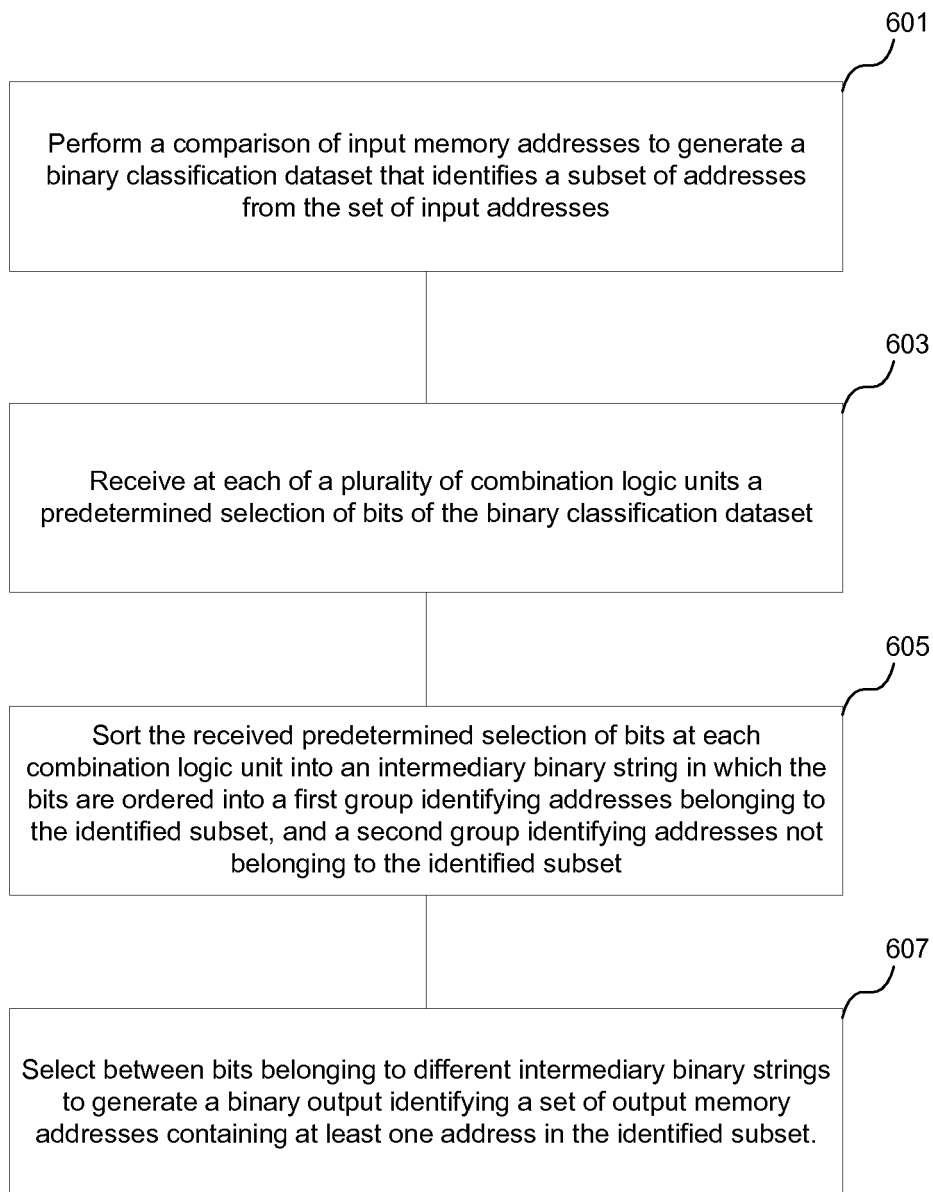


FIGURE 6




FIGURE 7




FIGURE 8



EUROPEAN SEARCH REPORT

Application Number
EP 18 19 4678

5

10

15

20

25

30

35

40

45

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (IPC)
A	US 6 684 301 B1 (MARTIN GREGOR J [US]) 27 January 2004 (2004-01-27) * abstract * * column 3, line 14 - column 6, line 16 * * figures *	1-15	INV. G06F12/06 G06F13/16
A	US 5 522 059 A (MARUSHIMA TOSHIKAZU [JP] ET AL) 28 May 1996 (1996-05-28) * abstract * * column 1, line 19 - column 5, line 7 * * figures *	1-15	
A	US 2006/129764 A1 (BELLOWS MARK D [US] ET AL) 15 June 2006 (2006-06-15) * abstract * * paragraph [0002] - paragraph [0022] * * figures *	1-15	
A	US 8 108 625 B1 (COON BRETT W [US] ET AL) 31 January 2012 (2012-01-31) * abstract * * column 2, line 7 - column 3, line 36 * * column 15, line 17 - line 35 * * column 17, line 46 - line 61 * * column 19, line 35 - line 40 * * figures *	1-15	TECHNICAL FIELDS SEARCHED (IPC) G06F
A	US 2008/183982 A1 (SUGAHARA TAKAHIKO [JP] ET AL) 31 July 2008 (2008-07-31) * abstract *	1-15	
A	US 6 810 470 B1 (WISEMAN JOHN R [US] ET AL) 26 October 2004 (2004-10-26) * abstract *	1-15	
A	US 6 438 660 B1 (REAMS BYRON L [US]) 20 August 2002 (2002-08-20) * abstract *	1-15	
The present search report has been drawn up for all claims			
Place of search The Hague		Date of completion of the search 23 January 2019	Examiner Mandato, Davide
CATEGORY OF CITED DOCUMENTS X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons & : member of the same patent family, corresponding document			

EPO FORM 1503 03.82 (P04C01)

50

55

**ANNEX TO THE EUROPEAN SEARCH REPORT
ON EUROPEAN PATENT APPLICATION NO.**

EP 18 19 4678

5

This annex lists the patent family members relating to the patent documents cited in the above-mentioned European search report.
The members are as contained in the European Patent Office EDP file on
The European Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

23-01-2019

10

15

20

25

30

35

40

45

50

55

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 6684301 B1	27-01-2004	NONE	
US 5522059 A	28-05-1996	JP 3010947 B2 JP H06161885 A US 5522059 A	21-02-2000 10-06-1994 28-05-1996
US 2006129764 A1	15-06-2006	CN 1786901 A US 2006129764 A1	14-06-2006 15-06-2006
US 8108625 B1	31-01-2012	NONE	
US 2008183982 A1	31-07-2008	JP 5137002 B2 JP 2008181399 A US 2008183982 A1 US 2013013887 A1	06-02-2013 07-08-2008 31-07-2008 10-01-2013
US 6810470 B1	26-10-2004	NONE	
US 6438660 B1	20-08-2002	NONE	